UK Computer Science Research: Vision and Opportunities

(Input to the EPSRC International Review Panel)

UK Computing Research Committee

Executive summary

1. Information and Communications Technologies (ICT) are crucial to progress; they can contribute both to wealth creation and to quality of life. (§ 1)
2. Computer science research is needed to feed the development of ICT where –though it is the most successful product of the twentieth century– many scientific and technological challenges remain. (§ 1.2)
3. Fundamental research has an intrinsic intellectual value; in computer science it can also have major long-term payoff. (§ 1.3)
4. UK academic computer science has a world-class scientific record and has successfully transferred many ideas to industry. (§ 2.2 and Appendix A)
5. The academic community has a vibrant and challenging research agenda (§ 3.1) focussed around
   - System design
   - Information powered society
   - Pervasive computing
   - New devices and Large Scale Integration on a Chip
   - Theory and foundations
6. Announced projects in "e-science" and "bioinformatics" present exciting challenges for some areas of computer science research. (§ 3.2, 3.3, Appendix B)
7. Our international competitors are investing more heavily in university computer science research than is the UK. (§ 2.3)
8. UK industry is investing little in computer science research. (§ 2.3)
9. The UK university system as a whole has experienced significant financial cutbacks over the last twenty years. (§ 2.3.1)
10. Plans show that UK research council investment in computer science research will not increase over the next three years; support could come from European Union programmes and there are significant EPSRC opportunities such as the Interdisciplinary Research Collaborations (§ 2.3.2) and the e-science (§ 3.2) and bioinformatics (§ 3.3) initiatives.
11. There are serious staff retention problems in UK academic computer science; these throw into doubt the future of many university departments. (§ 2.4)
12. To enhance scientific progress and maximise the potential impact on industry, a number of issues require urgent attention. (§ 4)
    - Long-term funding is required to establish and/or retain world-class research groups.
    - A way must be found to attract enough (UK and European) PhD students to feed through into both teaching and research positions in UK universities as well as satisfy industrial needs in the UK.
    - Recruitment and retention issues of academic staff must be addressed.
    - There should be a professional study of academic/industrial collaboration aimed at establishing a clear picture of the *long-term* impact of academic research as well as finding new modes of collaboration with industry for some aspects of computer science research.

Table of Contents

1 The case for computer science research

ICT (Information and Communications Technology) has a central and strategic role in modern society. Computer Science research is essential to support this role and to develop improved understanding of the foundations of computation.

1.1 ICT in modern society

Most workers in developed countries use information technology at some level. ICT is a driver of developed economies; furthermore, as society becomes more information driven, good systems offer competitive advantage.

The proportion of a developed country's GNP spent in building computer systems has become significant. This is not solely the cost of purchasing hardware and software; most large organisations employ computer professionals to build in-house systems. Because definitions vary, it is difficult to get consistent numbers, but figures such as two million programmers in the US (an annual cost of well over $200B) and an estimate that the UK is close to employing one million computer professionals, are sobering when the high cost and restricted availability of such professionals are considered.

Nor is there any indication of the growth in demand diminishing. In this century, advances in computing and communications technology will create a new infrastructure for business, scientific research, and social interaction. Developments such as wireless communication, the Internet, next-generation information systems, e-business, graphics and new media, are already transforming sectors such as telecommunications, education, medicine, finance, retail and entertainment, with immense potential for wealth creation and improved quality of life. New developments such as optoelectronics, nanotechnology and quantum computing pose fundamental computer science research challenges.

1.2 The future of ICT

In spite of their huge contribution, it would be a serious mistake to view computer systems today as ideal. Above all, there a continual hunger for new applications (this topic is addressed in subsequent sections). It is also apparent that many existing systems are neither reliable nor easy-to-use. These shortcomings both result from and are exacerbated by the wide use of computer systems. Developed countries are facing a shortage of computer professionals and in many cases are employing people with little fundamental computing education. The wide use of systems places demands for user-friendliness that were absent when the only contact between a company's mainframe computer and any non-computer professionals was printed output.

More economical and predictable ways of building computer systems could save each developed country sums measured in billions. Harder to quantify, but even more significant, are the savings that would accrue to the user community (which is far larger than that of computer professionals) from systems that opened up new application areas or enabled users to work more productively.

The observation that there is scope for improvement should not be taken as a criticism of the computer sector that has delivered so much already[1]. The hardware industry has achieved exponential (Moore's-law) progress in storage media and communications as well as processor performance. Software has also made spectacular advances with systems being built today that were unthinkable just a generation ago. There is also a continuing improvement in the productivity of creating advanced applications with ever more powerful

[1] The US PITAC Report [Ref 10] estimates that IT is responsible for 30% of economic expansion since 1995, and that it has kept inflation a percentage point lower than it would otherwise be.

tools and infrastructure being deployed to reduce the fraction of new code required. Neither hardware nor software designers can achieve exponential acceleration of the design process itself but both areas have learned to deploy tools and reuse more components as a way of building ever more powerful systems.

1.3 Types of research

It could be asked whether there is a need for more computing research; has the phenomenal progress of a half-century left us with an adequate pool of knowledge? It would indeed be unique for any subject as ubiquitous as computing to have achieved its pinnacle of knowledge in fifty years, and even a superficial investigation makes it clear that this is not the case.

There are many reasons for undertaking research. In the physical and mathematical sciences, support is given to fundamental research aiming at knowledge for its own sake. In computer science, there are goals that should be regarded in this category. But here even seemingly esoteric problems such as whether “P=NP” could have major and immediate implications for areas like security.

There is an obvious case for curiosity-driven research in any area as significant as computing. Understanding the fundamental notion of computation must rank as one of the key contributions of the last century and certainly played a part in the creation of computers. Equally important, although perhaps more elusive, is the fundamental notion of communication: much progress has been made and UK researchers are in the vanguard but more research is required. Computer science has also caused a re-evaluation of what constitutes a proof.

A more materialistic argument is the expectation of improved industrial competitiveness from computing research.[2] The past track record here is good (see discussion on the UK contribution in Section 2.2). There is no research agenda that can guarantee dramatically improved productivity in the creation of programs but research specifically targeting this (see Section 3.1) must offer a high potential return.

Another application-led case for research comes from the desire to improve the quality of life. Whether, for example, tackling medical problems or those of predicting effects on the environment, computing research can, in consort with other scientists, make a huge contribution to the wellbeing of everyone on the planet.

1.4 Research funding outside the UK

It is clear that the UK’s competitor nations are investing heavily in computing research. The US Government has recently increased its computer science research spend by around a third. Similar step increases have recently taken place in France, Germany and Ireland; Denmark has created an IT University. It would be useful to undertake a review of the percentages of the total research spending in various leading countries that goes on Computing Science but such a review must be conducted professionally since variations in definitions could otherwise result in meaningless comparisons.

2 The position in the UK

The overall research council spend through OST in the UK is £1.3B in the current financial year. EPSRCs share is £400M of which about £65M is the base line for the ITCS Programme

[2] It is important to keep the distinction between pure and engineering aspects of computer science in mind in the discussion throughout this report. It is a legitimate concern of governments to measure the industrial impact of research funding but this should not be the only basis on which research is assessed.

on grants and studentships.[3] It is encouraging that the recent trend has been to increase overall science spending but the current EPSRC Balance of Programmes exercise indicates that IT&CS will do no better than retain its planned level of funding for the next three years. There are however additional targeted programmes (see Sections 3.2 and 3.3).

There is a clear recognition that ICT is key to the UK's future. A recent study by the OST [Ref-1], the UK Government Department responsible for Science and Technology, suggested a target for the year 2005: with production of ICT goods and services adding 0.3% to the percentage GDP growth rate and responsible for 30% of UK employment growth. That report, and the recent white paper [Ref-2], regarded curiosity driven basic research as an essential step in achieving this goal, recommended the UK retain and strengthen its scientific and technological base within research-oriented universities, and stressed the need for "radical proposals" to replace conservative science.

Unfortunately, UK industrial research spending is at a lower level than its competitors but precise figures are difficult to obtain because of varying definitions of "R&D". It is however possible to obtain Europe-wide comparisons for RTD spending: as a percentage of OECD spend, the US figure is 52% where Europe is 17% and Japan is 22%. Thus the whole of Europe is spending less on Research, Technology and Development in IT areas than Japan.

2.1 UK computer industry

The UK is no longer a player in the mainframe or PC market. This obviously removes a whole avenue of potential collaboration for academics. In more specialised hardware, UK computing research has been key to major commercial developments. For example, companies such as ARM and Inmos have designed products that are sold in huge numbers and both companies acknowledge the key role of university research. In the software area, companies such as Autonomy, Virata, Baltimore and Sage[4] [5] have huge turnover.

Computing research has contributed to the development of high-tech clusters, notably in Cambridge where about 1000 high-tech companies generate £2B in revenue. Scotland manufactures 28% of Europe's PCs, and nearly 80% of Europe's workstations, with a rapidly growing microelectronics and systems-on-a-chip cluster around Glasgow and Edinburgh Universities, in collaboration with Cadence and Motorola. The M4 corridor is the home to many of the UK's software driven businesses.

Foreign companies, such as AT&T, Agilent, Cadence, Canon, Hewlett Packard, IBM, Microsoft, Nortel, Sharp, SRI and Xerox, have shown their confidence in UK computing standards with major research and development labs in the UK. While, in general, UK industrial R&D spend is low, UK corporations such as British Aerospace, British Telecomm, GEC, Logica, Marconi and Rolls Royce undertake significant computing research in-house and through university sponsorship.

2.2 Some examples of UK research

The United Kingdom has made crucial and seminal contributions to the development of computer science dating back to the theoretical underpinning for universal computation developed by Alan Turing in the 1930s, the world's first stored-program computer (1948); followed by sustained work on hardware architectures and key research results in theoretical computer science. Today there is world-class work in many sectors, a product of sustained

[3] Precise figures on grant *expenditure* are best obtained from EPSRC since groupings have changed.

[4] Sage is a little known success story: a Newcastle based company that has a market capitalisation of about £6B: its founder, Graham Wylie, wrote the first Sage accounting software while in Computing Science at Newcastle University.

[5] Company values are indicative and were taken circa Jan 2001

research programmes by research groups that have been formed over many years. Results range from fundamental theoretical insights to concepts and code transferred into industrial products. Just four areas are indicated here; for detail and a broader list of sectors, see Appendix A.

- **Theory, formal methods and foundations**: The UK is a world leader in the foundations of computing: Strachey and Hoare made crucial contributions to semantics; Milner's notion of mechanised proof was devised in the 1970s, developed through systems such as Gordon's HOL and Paulson's Isabelle and is only now being exploited by Intel and others in production hardware verification. Plotkin and Abramsky are internationally recognised for their work in semantics and Jerrum for his study of rapidly mixing Markov chains (key to modelling processes such as network routing).
- **Programming languages**: Edinburgh's Standard ML has been enormously influential, combining the discipline of strong typing with the flexibility of polymorphic type inference (themes further developed by Haskell, in which Glasgow provided major players). As so often, tracing the impact is made more difficult by name changes: the polymorphism present in Haskell turns up as "generics" in C#. (Another technique pioneered for functional languages is garbage collection that is used in Java.) Milner's π–calculus is the foundation of techniques aimed at giving a sound basis for 'mobile' and 'global' computation, notably current work on Ambients by Cardelli and Gordon at Microsoft, Cambridge.
- **Dependability**: Brian Randell has led the Newcastle team since the seminal architectural concept for coping with design errors: "recovery blocks" provide a way of making alternative algorithms available at run time (more recent developments are sketched in Appendix A). This avenue has led to a transaction system marketed by a spin-off company (Arjuna Solutions Ltd. [Ref-14]) which was recently acquired by HP. Piper and Mitchell head a noted security group at RHBNC which spawned the security company Baltimore, with a market capitalisation of £3B.
- **Computer architecture**: The UK is still a world leader in specialist hardware design: the Transputer development is a model of industry/university interaction in technical innovation, and is in 40M set-top boxes (60% of the world market). Processors designed by ARM Ltd (market capitalisation of £5B) shipped 400M in year 2000. These are the volume RISC standard in markets such as portable communications, hand-held computing and embedded devices: they are based on work by Furber, now at Manchester University, who is working with ARM on AMULET, a low-powered asynchronous ARM processor.

2.3 Research funding

The bulk of UK computing research is carried out in Universities: academic staff combine teaching and research, and permanent research posts are extremely rare. Full time "research associates" are employed on short-term contracts funded by specific grants from either UK research councils or the EU. About half of all research funding comes from the institutions' Government grant-in-aid, the rest from project funds awarded competitively by EPSRC, and other (mainly industrial) sources. The EU-funded ESPRIT programme, now IST-5, has provided long-term investment in a variety of industry/academic basic research projects.

2.3.1 Direct university funding

Governmental funding to UK universities is calculated on two parameters: their teaching activity and their research activity. These interact, since most institutions; departments and individuals are involved in both activities.

As successive UK Governments have followed a course of expanding the percentage of the population that receives tertiary education, they have significantly reduced the resources per student. Fortunately, the *nominal* per capita fee for computer science students has remained stable in real terms. The *effective* per capita fee for these students has however dropped severely due to the rapid growth in computer science student numbers and the tendency of institutions to cover unprofitable activities by taxing computer science. Only a proportion of new students result in new funded places because resource distribution inevitably lags behind growth. In most cases this results in a subsidy from research to teaching.

The research component of Government funding is based on two parameters: the level of research activity (e.g. research active staff, research students, etc.) and the quality of the research based on outputs (papers, patents, etc.). These two parameters are determined by a quinquennial Research Assessment Exercise (RAE); the current round completes in the summer of 2001[6]. However, the lack of responsiveness in teaching funding means that research has suffered in consequence, resulting in unbalanced cross-subsidies and leaving many researchers unable to allocate adequate time for research.

2.3.2 Research councils

The prime funding for computer science research in the UK comes from EPSRC but their forward look shows IT&CS funding as not increasing. There are however some encouraging developments. In year 2000, five *Interdisciplinary Research Collaborations* (IRC) were funded by EPSRC (and MRC).[7] These large distributed interdisciplinary teams have been funded as six-year programmes of long-term research: altogether 20 Universities and around 45 companies are involved. The areas of research are

- **Advanced knowledge management** Lead site Southampton, Nigel Shadbolt: improving knowledge management to prevent information overload.
- **Dependability of computer-based systems** Lead site Newcastle, Cliff Jones: improving the safety, security and reliability of complex systems that involve people, computers and communications.
- **Photonics innovation** Lead site St Andrews, Wilson Sibbett: ultra fast communications using optoelectronics above terabit speeds.
- **Equator: technological innovation in physical and digital life** Lead site (now) Nottingham, Tom Rodden: embedding computers into everyday objects and environments.
- **Medical informatics** Lead site Oxford, Mike Brady: transforming medical images and data into usable clinical information.

The total funding for these projects is more than £40M spread over 6 years.

2.3.3 Faraday partnerships

Faraday partnerships are a novel attempt at helping technology transfer. To quote from the Web announcement[8]

> Over the next four years the Faraday Partnership projects will be supported by £4.8M from the DTI and £4M from the Engineering and Physical Sciences Research Council (EPSRC). Faraday Partnerships enable firms to access high quality research and expertise of the industrial research organisations. They employ "technology translators" - people who can act as a link between research and business, ensuring

---

[6] In the 1996 RAE, 90 academic institutions listed research active staff in Computer Science; in comparison only 56 institutions returned staff under Physics and 62 under Chemistry.

[7] Further reading: EPSRC Newsline Issue16 on IRCs
www.epsrc.ac.uk/Documents/About_EPSRC/Corporate_Publications/Newsline_Journal/Newslin16/digital.htm

[8] www.smithinst.ac.uk/PressRelease.html

that good ideas are properly developed and that researchers understand the potential that their work offers industry.

2.3.4 EU

The European Union has funded computing research since the first ESPRIT programme. A very useful step was the recognition of *Basic Research Actions*.

UK researchers fare well in EU funding. In the current (IST-5) programme, UK co-ordinated projects represent about 13% of the budget and in some areas are over 20%.

2.4 Staffing

The attraction and retention of people is now at crisis level in several key areas of computer science, as reported in [Ref-5]. We are in an international marketplace for talented individuals, and competition is acute, not just for star researchers, but also for the post-docs and research students. We are losing staff to industry and abroad; for example, one top-rated University Department recently lost three Professors to industry, two within the UK and one to the USA. A major factor is uniform university salaries[9] (regardless of discipline), a young Computer Science lecturer will be on a salary of about £23K. Moreover, a PhD student will be on a stipend of less than £9K. Starting salaries for computer science students with only a first degree will often exceed those of lecturers; moreover, public service salary scale increases are controlled by central government and have for decades risen far slower than average industrial increases. It is increasingly difficult to attract young people into academic research and we are in danger of failing to hire the next generation of researchers and lecturers. In fact, it is an indication of how good the best departments of computing are that they do retain top-class researchers in spite of the obvious attractions externally. But the dependence for the future of the subject and the supply of research-trained staff to industry is so great that something must be done to improve recruitment and retention of staff.

3 Research challenges

This section sets out a strategic view for research in Computer Science. We first list five strategic scientific goals that incidentally indicate why bioinformatics and e-science are seen as timely challenges.

Two current major OST/EPSRC interdisciplinary initiatives, in e-science and bioinformatics, will provide a framework for interaction between computer science and other sciences. Both offer potential for extending and applying core computer science concepts, and for an impact on scientific practice, as well as a laboratory for computer scientists to test out their theories and methods on complex computer systems and massive, highly structured, data.

3.1 Five key research challenges

We identify the following five themes as key research challenges.[10] Each builds on established UK strengths and, taken together, they offer the opportunity for long-term scientific advance, technological innovation and commercial exploitation, in areas which are ripe for significant progress and where the UK is already strong. Each area deserves sustained, well-funded research programmes and significant investment in tools and personnel. A robust infrastructure is also necessary, to pioneer new techniques and comprehensive experimentation, evaluation and assessment.

---

[9] An exception has only been made for medicine where starting salaries for academics might be a third higher; the non-clinical (full) Professorial minimum is £37K whereas for clinical it is £49K.

[10] Other interesting attempts to determine a research agenda are given in [Ref-12, Ref-13].

3.1.1 System design

The total cost of ownership of poorly designed computer systems dwarfs the costs of the UK's one million programmers; progress in system design is therefore essential. If hardware and software were the most significant achievement of the last century, better ways of designing complex, dependable and usable systems are crucial for this century.

Our society, industry, commerce, services (such as healthcare) and Government depend today on a huge, ever more complex network of distributed software systems, lashed together from massive numbers of software components using a Babel of languages and protocols. The quality of the individual components is questionable and the overall behaviour of such systems, especially as to how they can be maintained, managed, changed, improved, extended and made secure is poorly understood.

Today's huge computer systems are the most complex creations of human minds and it should be no surprise that adequate design methods require research and investment. The demand for software has grown faster than our ability to produce it in a reliable way, and we software is needed that is more usable, reliable, and powerful than that which is being produced today. Particular challenges are: making software easier, quicker and cheaper to design, create, maintain, and inter-operate; designing and building very large computer and information systems, systems that evolve; and extending the ways humans interact with computers, for example by voice or touch.

Progress will capitalise in part on existing strengths in distributed systems, HCI, artificial intelligence and cognitive science, information engineering, software engineering, dependability, programming languages, and theory, formal methods and foundations. Progress requires fundamental research in: accelerated software development methods for robust reusable software; creation of evolvable software which can cope with unpredicted contextual change; new and improved forms of human-computer interfaces and interaction; capturing, describing, managing, analysing, integrating and explaining information; evaluating solutions, methods and technologies**,** "useful" complexity measures for performance and cost.

Although hardware performance has improved enormously, there are also applications where performance is still critical. These might be in a minority as a proportion of programs but together probably consume a significant percentage of total machine cycles.

3.1.2 Information powered society

Almost every aspect of our society depends on computer-supported information processes and deserves increased research investment. The ability to process, move and store data has grown by a factor of more than a million in 30 years and such growth is expected to continue for the next 30 years. The UK is already a world leader in the production of high quality digital content (for example, the BBC's web site receives the highest number of hits of any content provider in Europe, Soho net is a world leader in delivering material to the film industry, a major part of Reuters' information systems are developed in London where they were founded, and the Open University –whose courses have been followed by over two million students– is a major provider of on-line education).

The UK's growing wealth of data has the potential to power knowledge discovery, stimulate creativity, improve life-long learning and ensure that decisions are well informed. But this potential is hard to realise without further investment. Challenges arise at every stage from information gathering, through information integration to information extraction and presentation. Two typical issues are:

- how to describe the semantics, quality and provenance of data sufficiently well that it can be *accurately* and *automatically* integrated with other data without inhibiting the variety, creativity and ingenuity of the information collection or generation processes; and

- how to encourage the *accurate* retrieval and presentation of information from the gamut of data sources using the full range of techniques from formal queries, through data mining and information retrieval to guided browsing.

To realise the potential of ubiquitous varied and evolving digital-information flows, significant extension of theory and practice is needed. In the UK this will build on and combine existing strengths in groups active in HPC, graphics and digital media, HCI, artificial intelligence and cognitive science, information engineering, software engineering practice and theory (see Appendix A for details).

3.1.3 Pervasive computing

Billions of mobile, intermittently connected, low-powered computers are expected as "active badges" and "condition monitors" on virtually every domestic and industrial appliance, component and commodity. A consequence is that computing will become pervasive, with networks supporting a wide range of nodes from very powerful servers, through context-aware palm-tops, to intelligent devices and smart fabrics. Companies, such as IBM and Sun, already include pervasive computing in their target markets, and use it in their assembly plants to track components and supply versions of Java designed for these applications.

New network and application-software architectures will be needed for the radically changed distributions of intelligence, the variety of computation contexts and the dramatically expanded number of communicating agents. With the UK's traditional strengths in networks and distributed systems, we are ready to be major players.

A key issue is programmability, that is programming the network, resources, services and applications, and opening up the networks to third parties. With such a plethora of computational components, virtually everything must be automated and parameterised by device, context and policy descriptions. Users should simply specify policy and function. A network must continue to deliver appropriate quality of service to all users, and be robust and resistant to technology faults, human error and malicious attacks and accommodate a rich variety of data and protocols. To achieve this goal requires fundamental research in: system architectures, especially for dynamic run-time reconfiguration, open standards and APIs; Quality of Service, at the application level, and incorporating mobility; active networks, and context and network aware applications, devices and interfaces; performance management and models of resource utilisation, costing, and behaviour; protocols, languages and standards for interoperability of services and applications. Progress will draw on existing strengths in architectures, networks and distributed systems, HCI, artificial intelligence, cognitive science, information engineering; software engineering, dependability, programming languages and theory.

3.1.4 New devices and Large Scale Integration on a chip

Research in industry and academia is delivering significant advances at the physical layer. For example, nanotechnology combined with digital electronics is yielding a succession of ever faster, smaller and more versatile versions of the laboratory on a chip (LoC). Combined on chip with photonics the LoC can accomplish spectrographic analysis. On-chip radio can receive and transmit messages. The advent of black silicon suggests rapid advances in optical sensing devices. Nanoengineering has developed pipettes, motors, fans and other micro-actuators. These are examples of a rapid expansion in the ways in which a complete system on a chip can interact with its environment. Such systems will be deployed as cheap, mobile and semi-autonomous sensors, for example, ingestible diagnostic devices, environmental pollution sensors and personal health monitors.

The challenges for computer science arising from these advances in integration, chemical and physical interaction, and miniaturisation are legion. The design of these systems raises new challenges for CAD, where many design components have to be integrated without inconsistency by engineers from several disciplines who cannot be expected to comprehend their detailed operation. The operating systems and communication protocols for these devices will raise new safety issues to be addressed in the context of scarce resources, such as bandwidth and power. The applications using these devices have to reduce the cost of installation and update to a minute fraction of today's costs. Swarms of such instruments will be integrated into larger systems that must be dependable and updateable despite the frequent, often permanent, loss of service from individual instruments. New challenges in observing and understanding human behaviour and sharing responsibility arise. For example, suppose that farmers use LoC palmtops to test minute samples from their stock and people use them as implants to monitor their own and other people's health. Treatment decisions based on readout may save veterinary and medical labour, but how can we be sure they are safe? How would we feed epidemiological data to the remote devices, e.g. when there is an outbreak of food poisoning? How would we spot a recurrent error or detect a drift in a reusable device's calibration? Such questions raise architectural design issues about divisions of responsibility and supporting infrastructure that require computing science research and insights.

These advances in devices will affect every branch of computation. A laptop or palmtop may not only detect smoke and warn its user, it will use other sensors to suggest the best escape route, or raise an alarm when it detects anomalies in its user's pulse. We cannot begin to imagine the full potential of more powerful mobile communicating, context-aware computers, with a growing range of sensory devices and actuators.

We are well placed to capitalise on the UK's existing strengths in optoelectronics and nanotechnology: progress will build on and extend our strengths in hardware, communications, HCI, cognitive science, software engineering and dependability.

3.1.5 Theory and foundations

The UK is a world leader in the logical and mathematical foundations of computer science, and dynamical systems for alternative models of computational phenomena, both natural and artificial. Long-term research (see Appendix A for details) is only now manifesting its far-reaching practical consequences in explaining network behaviour, allowing accurate modelling of complex distributed phenomena, and understanding speech and language. Future challenges include developing models, theories and implementations using experimental, logical, statistical and mathematical techniques, so that we can understand, simulate and predict the behaviour of massive data, massive networks and massive computations in natural and man-made computer systems.

For example, what's the right way to model massive, any place, any-time computation using techniques from logic, probability, statistical mechanics or dynamical systems? How can these models be used to design networks with desirable security or traffic flow properties, and what theories and properties will we need to study in the future? In artificial intelligence research, how do we unify probabilistic and machine learning-based methods with high-level symbolic representation and reasoning, or develop models that unify cognitive function and inference? How do we turn our theories into workable metrics or visualisation techniques to address practical problems?

Progress requires fundamental research in: representation and reasoning; logic and theoretical computer science; pure and applied mathematics; probability and stochastic techniques; simulation and modelling. It will build on and extend our existing strengths in artificial intelligence and theory, formal methods and foundations, and impact each of our key areas to a significant extent.

3.2 The e-science initiative[11]

The e-science initiative, announced in the OST spending round of Summer 2000, represents a Grand Challenge for the UK's science and engineering base and most of the research strategy set out in Section 3.1 finds a ready application in e-science. Science is a knowledge-based discipline and, consequently, information –its storage, retrieval, exchange and manipulation– features heavily in the challenges presented by the scientific community. Science is naturally

[11] Further information: Appendix B http://www.research-councils.ac.uk/ostinformatics

distributed and is now conducted on a global scale. The data, and hence information, streaming from high throughput experiments is growing exponentially, and the varieties of information that must be combined is expanding, is related in complex and ill-understood ways and comes in many representational forms that are not immediately accessible by computing techniques.

This requires “the GRID”: a fast, reliable, distributed and interoperable computing network, together with knowledge infrastructure and support for the scientific process. There are numerous exciting opportunities for computing research: tools and systems to handle all aspects of organising and monitoring computation, middleware for information access and integration, and systems for data mining and knowledge management.

Particular impact is expected in bioinformatics, medical science, physics (e.g. virtual observatories, data and computing grids for the particle physicists’ LHC experiment) and meteorology (e.g. global warming studies and medium-range weather forecasting). The development of e-science and the GRID, particularly for subjects such as biology, ecology and oceanography, which are multi-faceted, should give UK researchers an opportunity to explore such issues. But this depends on investment to encourage UK computing science researchers’ involvement throughout that development. Other platforms providing high volumes of data and rich variety include medical informatics. Research using such platforms is crucial for the future of e-business in the UK as these platforms combine two key properties: they are accessible to academic researchers; and their complexity, variety and rate of evolution, emulates the more demanding aspects of e-business.

3.3 Bioinformatics and life sciences initiative[12]

The sudden wealth of data in biology and medicine portends a bonanza of discoveries advancing our understanding of life processes, but it demands extensive and rapid growth in new and traditional forms of computation. This will be a powerful driver for computing research and an opportunity to deploy the fruits of the research agenda set out in Section 3.1. One such approach is the facilitation of co-operation among scientists enabled by networked computation. Bioinformatics offers challenges and opportunities as biomedical research depends on much global and interdisciplinary collaboration and uses a rich diversity of methods and data.

These EPSRC-BBSRC initiatives have funded a variety of innovative projects, together with four chairs in Bioinformatics at Manchester, UCL, Imperial and Oxford. The UK is currently a world leader in collaborative bioinformatics (see Appendix A for details). Bioinformatics is a young discipline – the collaborative aspect is a consequence of the confluence of the Human Genome Project, the design of novel approximate matching algorithms and the distributed information infrastructure provided by the Web. Accumulated knowledge and powerful computational techniques increasingly suggest discoveries that are borne out by experiment, and this is expected to become more prevalent.

The first phase of bioinformatics, genome sequencing, is well advanced. The relationship between bioinformatics and computer science here applied established computing techniques: pattern matching, databases, process management, etc. Most computation was localised, consequently it could be achieved by labour-intensive integration techniques. The future requires significant innovation. The variety and rapid development of the ever-growing repositories of knowledge about biological systems from proteins and genomes to ecologies and evolution portends great potential for discovery. But this will be widely accessible only if knowledge integration and much-improved indexing, searching and approximate pattern matching techniques are developed. As the biomedical knowledge base grows, researchers need to locate and transfer knowledge between contexts, exploiting nature’s conservatism. As our understanding of life’s processes, such as the functions of genes, develop, their models become more computationally challenging while the potential results of simulation become

[12] Further information: http://www.bbsrc.ac.uk/science/initiatives/bioinformatics.html

more directly applicable. An example of the response to this challenge via new architectures and algorithms is IBM's Blue Gene project.

4 The way forward

Our goal is ambitious: world leadership in the five key research areas listed in Section 3.1, successful long-term exploitation and universities as vibrant centres of international computing research. The quality of our Computer Science base and our industry/university interaction will be key to capitalising on the wealth of opportunities.

To achieve this impact we need long-term research programmes with freedom to cover multiple aspects of big problems, and resources to do the basic science, run experiments and field a variety of interactions with industry. We need facilities to support large and long-running experiments evaluating such issues as the effectiveness of design methods, software architectures, evolution mechanisms, distributed system technologies and automated integration mechanisms under realistic conditions.

As we have seen in the cases of ARM, Autonomy, Baltimore and Virata, long-term UK Computer Science research programmes have had significant industrial impact. No less is expected of our new IRCs and the interdisciplinary work in bioinformatics and e-science; all aim at significant industrial interaction.

We should not be complacent: the UK has repeatedly failed to capitalise fully on its science base, and the reasons, including an anti-technology culture, and industrial under-investment in R&D, run deep (see Appendix C for a summary taken from a recent OST study [Ref-8]). New modes of interaction between industry and academia must be sought.

4.1 Improving the environment for Computer Science research

At their best our university computing departments are exciting places. They are driven not by monetary reward or management incentives, but by the energy of creative thinkers who are passionate about Computer Science and its potential. Unfortunately battles for scarce resources and global competition for the brightest researchers frequently erode this ideal.

In order to achieve critical mass in a number of world-class laboratories we need:

- support for large, long-term research programmes, encouraging and sustaining strong centres and communities to work on big questions and providing the infrastructure they need to support basic research. This should include large, long-lived experiments and a variety of modes of collaboration with industry;
- a review of mechanisms for university-industrial collaboration in the widest sense, to provide flexibility and establish industrial interaction as a third force to complement teaching and research;
- support for and flexibility in overseas interaction, particularly through involvement in NSF and other US projects, to complement existing strong links with EU scientists;
- support for mechanisms that allow Computer Science researchers and academics to be rewarded at levels that are comparable with the UK's international competitors.

This last point is of particular importance because of the low-level of academic rewards in the UK. For example, a UK research student receives less than £9K pa, compared with £25K or more in Europe; an experienced UK professor's salary is 50%-60% of his or her German equivalent. Furthermore, a rapid increase in demand for Computer Science undergraduate courses has led to a situation where Computer Science academics have faced increasing workloads but with lower resources per student. The lower resource level has arisen as subjects with decreasing student numbers are cross subsidised by income from computer science, and because of the partial and delayed response in funding from the Higher Education Funding Councils

Government White Papers [Ref-1] indicate both increasing growth in student numbers and a commitment to high-quality, adventurous research. However, unless there is a direct injection of resources into departments where these activities have to be sustained it will be impossible to achieve both of these aims.

4.2 Establishing a policy framework for Computer Science research

Science policy and funding is the business of the Office of Science and Technology. OST funds computer science research primarily through the EPSRC. Government maintains its own knowledge base and advisory structures: such as this International Review, and other mechanisms such as EPSRC's TOP and UP panels, and the Higher Education Funding Councils' (HEFCE) Research Assessment Exercise (RAE).

Policy makers can help us in creating the best possible environment for research, by:

- taking the lead in recognising the importance of long-term scientific research in technological issues, and its relationships with shorter-term R&D in industry;
- working closely with academic and industrial leaders, and the NAS CSTB study of computer science, to devise appropriate criteria, metrics and structures for use in policy making.

From the UKCRC side, we are prepared to offer strategic advice (cf. Section 3.1) and revise the strategy in the light of changes in need and potential.

The relationship between research in universities and in industry is central to policy making. While there must always be a place for curiosity-driven research with no obvious industrial application, it is right to expect the engineering aspects of any discipline to focus the majority of its research on areas that have long-term industrial potential.

A long-term perspective is essential for both research and technology transfer. Research ideas often take time to become effective: for example object-oriented programming and design took off on an industrial scale more than twenty years after the seminal contributions of Dahl and Nygaard. Similarly, Microsoft bought a company providing extended static checking thirty years after the concepts were described in Hoare's seminal "axiomatic basis" paper.

For many years, science policy in both the UK and Europe favoured research with shorter-term objectives and we are convinced that this has undermined university/industry technology transfer. In essence, the short-term industrial objectives changed as a result of changing technology, new products etc. thus making the research irrelevant. We strongly believe that it is only by taking a long-term view where research objectives may be more stable that we will be able to make a significant impact.

A major difficulty that policy-makers have faced is the lack of trustworthy performance indicators that allow the Computer Science research in the UK to be compared with comparable research elsewhere. Bibliometric data is particularly unreliable because of the different research cultures in the US and the rest of the world and because of the rapid rate of change of the discipline. World class research in new areas (such as ubiquitous computing) may not appear in bibliometric data because of the delays that are inherent in constructing these data sets.

The research community, policy makers and industry must work together to establish criteria that allow the scientific and industrial impact of Computer Science research to be established. These have to consider not just excellence but also timeliness, connectivity and adventure and must look at the impact on local as well as national and multi-national companies. A starting point for this might be the US Computer Science and Telecommunications Board study [Ref-6] to identify key questions and generate better understanding among scientists and policy makers. NSF is commissioning a study of Software Engineering [Ref-7]. Again, the UKCRC would be pleased to assist in developing these criteria.

4.3 Improving interaction with industry

Effective interaction with industry is critical to the success of much computer science research. There have been failings in this area although these have not been confined to academics. Indeed, having given the UK an unprecedented start in the hardware industry in the 1950s and 1960s, academics might doubt that the main problem with exploitation of world-class research comes from their side. However, the picture is not as bad as is sometimes painted (e.g. in studies such as [Ref-5] that are methodologically seriously flawed). The variety of successful interactions between universities and industry are often undocumented by formal processes. Tracing impact is made more complicated by the change of names of research ideas as they are embodied into products.

Industry clearly faces problems in interacting with researchers and these are particularly acute in Computer Science because of the rapid pace of advance of the discipline itself and of the business environment where computing technology is deployed. We have all been involved in research projects with industry that started well but where interaction petered out because of changing business priorities. Interaction with universities is particularly difficult for small and medium-sized businesses without in-house research facilities yet it is those businesses that have the responsiveness and the potential to benefit most from applying new research.

Whatever the situation has been in the past, it is pointless to ascribe blame. Rather, we must now consider how academia and industry can co-operate in developing the Computer Science discipline to their mutual benefit.

We are convinced that the principal difficulties in transferring technology from universities to industry have arisen for the following reasons:

1. Academics have insufficient time to devote to building relationships with industry and to develop a deep understanding of the real problems and needs of industry. This has partly arisen because of demands of teaching and administration but the approach to funding is also a contributor – the best researchers spend too much time writing (small) research proposals to maintain their existing research groups. This brings a concomitant high volume of peer reviewing the standards of which are consequently more difficult to maintain.

2. There is a critical ‘gap’ between the type of research funded by the Research Councils and the research results that are required to convince industry that a strand of research has the potential for industrial exploitation. Essentially, research council funding allows the development of preliminary prototype systems to demonstrate research results but does not provide funding to demonstrate the engineering viability of these results. Because of prevailing pressures, industry cannot take the risk of adopting preliminary results and, consequently, potentially valuable research is unused.

3. There is a lack of real data and testbeds from industry that can be used to validate research ideas. Understandably, industry has been concerned to maintain confidentiality and to avoid incurring the costs involved in making operational systems available especially when the research is still at a preliminary stage.

4. The fact that the UK computer industry is in many areas far smaller than that in the US limits collaboration.

5. The UK still does not have an entrepreneurial culture. Although the situation is improving, the rapid transfer of technology through university start-up companies is still relatively uncommon. The attitude to the ownership and management of intellectual property that is adopted by some universities may exacerbate the problems of promoting start-up businesses.

Without doubt, academics must take action to improve the transfer of technology to industry. However, we believe that there are a number of steps that the EPSRC might take to promote university/industry co-operation and which help address the above problems.

Firstly, the EPSRC can contribute to reducing the overheads involved in maintaining teams of world-class researchers. We welcome initiatives such as Interdisciplinary Research Collaborations and 'platform grants' that focus on long-term team building and would encourage the Council to continue and expand such initiatives. We would also support closer links between research funding from the EPSRC and HEFCE so that longer-term research can be supported. However, we understand the difficulties of achieving this.

Secondly, and critically, we believe that all participants would benefit from the funding of large-scale demonstrator projects in universities that can serve as both a vehicle for exploring the problems of practically exploiting research ideas and as a testbed for trying out new ideas. These long-term experiments (which are common in other disciplines such as particle physics) would serve as a focus for PhD projects and replicate, to some extent, the practical industrial conditions that allow researchers to assess and evaluate their ideas. Practical demonstration in such an environment will be the most effective way to convince industry to commit further funding and take up the ideas for exploitation.

Thirdly, the EPSRC is uniquely placed to bring together researchers in universities and 'problem-owners' in industry. All too often, contact between industry and universities is accidental and we are convinced that many fruitful opportunities for collaboration are missed. By proactively supporting activities such as industrial fellowships, workshops and demonstrations, the EPSRC can help reduce the barriers between industry and universities.

Fourthly, we believe that the current model of PhD funding is flawed. The 4-year limit before penalties are applied to universities does not allow PhD students to explore the industrial potential of their research or, indeed, to spend a period of time working in industry during their PhD. It encourages 'packaged PhDs' that can clearly be completed in the time available and discourages universities from steering students towards long-term, risky projects. We are not suggesting that there should be no limit but that an extension of this limit to 5 or even 6 years would allow PhD students to be able to work more closely with industry and to form start-ups to exploit their research. (This might be conditional on demonstrated industrial interaction, e.g. at least three months spent working as an intern in a collaborating company.)

Finally, we are convinced that an independent view on this matter would be of benefit to all parties. One possible way forward would appear to be for OST to fund an independent (SPRU or PREST) study of industry/academic co-operation. This should focus bothF on getting the facts and making recommendations for better two-way collaboration. Naturally, the members of the UKCRC will be delighted to assist with such a study.

Appendices

A Landscape of UK research strengths

The descriptions in this appendix should be taken as *indicative*, not definitive: we intend to give an idea of UK strengths in each area and some future directions, there is no attempt to list everybody active in each area nor all of the results and goals whether scientific or technology transfer.

Hardware and architecture

UK academic research has made a formidable contribution to the evolution of machine hardware and architectures. Many academic prototype machines have led directly to commercial manufacture, and ideas as important as virtual memory have their origins in UK universities. The timescales for this work are quite long — the research that has most visible commercial impact today was carried out in the early 1980s (the development of the ARM, the transputer and networking).

Manchester (Williams and Kilburn) and Cambridge (Wilkes – a Kyoto Laureate) produced the world's first electronic stored program computers; Manchester (Kilburn) went on to design four further generations of machine which transferred more or less directly to commercial products; when the Atlas was produced, it was the most powerful computer in the world; in the early 1970s work on MU5 (Edwards et al.) led to architectural innovations such as branch prediction and to the development of the ICL 2900 series, the mainstay of ICL's mainframe product range to this day. Research at Imperial (Darlington et al.) and Manchester (Watson and Watson) led to the ICL Goldrush.

In the late 1970s, work at Warwick (May) and Queen's, Belfast (Hoare) led on to the creation of the transputer and Inmos in the 1980s. The transputer is still manufactured (by STMicroelectronics) in volumes of around two million per year. Also in the 1980s Manchester (Gurd and Watson) had a leading role in the development of dataflow machines, an approach that has lost favour as the basis of large-scale parallel machines but still influences the design of out-of-order issue units in superscalar processors.

Through the 1990s to the present-day, world-leading research into asynchronous systems-on-chip continues at Manchester in close collaboration with ARM. ARM Ltd has a market capitalisation of £5B and in 2000 shipped 400M ARM processors, which are the volume RISC standard in markets such as portable communications, hand held computing and embedded devices. They are based on work of Steve Furber, now at Manchester, who is working with ARM on AMULET, a low-powered asynchronous ARM processor. AMULET is on the verge of its first commercial exploitation.

The Alba Centre is an ambitious project of Scottish Enterprise and four Scottish Universities, with inward investment from Cadence and Motorola, for R&D of system level integration for systems on a chip: current projects include measurement tools for environmental and medical applications.

The future challenges in hardware and architecture, at least in the short to medium term, include addressing the need for pervasive systems, many of which are mobile, and all of which are interconnected. They must support new and demanding data types, particularly for real-time multimedia and signal processing data streams. Power-efficiency and cost issues will continue the drive towards system-on-chip solutions with the concomitant emphasis on design automation, chip-area networks and design reuse.

In the medium to long term (beyond this decade) the physical limits of CMOS technology will have a major impact, and further progress in performance and power-efficiency will require new technologies. Various quantum devices are positioned to carry computing beyond where CMOS can go, but these will principally affect the lower levels of design (clocking strategies and on-chip communication schemes, for example). Beyond quantum devices is the prospect of quantum computing - exploiting quantum superposition in ways that present completely new computing paradigms.

Near future research in this area should reflect a balance between the relatively certain prospects of systems-on-chip, mobility, ubiquity, and extended data type support, and the more speculative prospects of quantum devices and quantum computing.

Communications networks and distributed systems

UK researchers have made significant contributions to the interconnected fields of object-oriented, distributed systems, multimedia and networked systems. Many results have been developed through close collaboration with industry, in some cases leading to commercial spin offs.

In computer networking, work at Cambridge (Hopper) led to the Acorn Econet, an HDLC technology that was the mainstay of educational computing in the UK in the 1980s. The Cambridge Fast Ring was an early example of an ATM network, it was the first high speed digital network platform. In the early 1980s, the Universe project, led by Cambridge, QMWC, Imperial and RAL, connected together local area networks using satellite links. It was an early example of mounting a larger collaborative engineering project.

The 50Mbit/second local area network was designed and implemented in the Cambridge Computer Laboratory in the mid 1980's, using a slotted ring structure, a development of the earlier Cambridge Ring. The Laboratory subsequently became a world leader in ATM switching networks, carrying out the first detailed traffic measurements and ATM pricing (Leslie). The Laboratory, in conjunction with Glasgow University (McCauley) and through the Pegasus I and II Esprit projects, developed the novel Nemesis operating system for distributed multi-media. Its spin off, Nemesys Research Limited, exploited the results, producing hardware and software to allow cost effective video connection to ATM networks. (Nemesys was acquired in 1996 by FORE Systems who were in turn acquired by Marconi in 1999). Virata, now a Silicon Valley based company with market cap of $3B, was another spin off from ADSL work at Cambridge.

In the mid 1980s, the ANSA project, began as an Alvey project and continued with EU funding. It involved an extensive consortium of industry and academia. The project's research on distributed platforms made a significant impact on CORBA.

One of the first distributed fault-tolerant systems was developed by Shrivastava at Newcastle, leading to the highly successful Arjuna system — object-oriented fault-tolerant distributed environments for workstations. These results subsequently led to the product OTSArjuna, a C++ implementation of the Object Transaction Service (OTS) specification from the OMG, marketed through the spin off company Arjuna Solutions.

Lancaster (Blair, Hutchison, Shepherd) develops distributed multi-media systems based on ATM networks, and the group has made significant contributions in the areas of quality of service management; multimedia communications services and protocols; multimedia information storage, mobile multimedia computing and CSCW. At QMWC, Coulouris, Dollimore and Kindberg wrote one of the seminal texts in distributed object-based systems (*Distributed Systems – Concepts and Design*) and develop support for distributed groupware, and real-time multi-media applications.

UK researchers play a major role in network management in large scale distributed systems. The first public-key authentication protocol using a key distribution centre was developed in the UK in 1978 (Needham and Schroder) in collaboration with Digital's SRC in the USA. More recently, heterogeneous operating systems and global programming systems are being developed (Connor, Glasgow and Strathclyde; Dearle, Stirling and St. Andrews); novel theories and implementations of distributed intelligent agents, policy-based manager agents, and configuration languages and support environments have been developed at Imperial College (Kramer, Sloman). Scottish academics play a leading role in the area of interoperability within communication services (Magill, Strathclyde and Stirling, Turner, Stirling; Calder, Glasgow).

Pervasive computing provides key future challenges, including: managing heterogeneous networks with more than a million nodes, delivering a demanding range of services economically, and simplifying the task of building and maintaining software running over these complex networks.

High Performance Computing

UK computer science has been heavily involved in all aspects of high performance computing (HPC) ranging from the development of formal notations to define languages and systems through to the construction of tools and applications. Of particular significance has been the work and development in the area of parallel computation. The pioneering work of Hoare and Milner with notations like communicating sequential processes (CSP) and a calculus for communicating systems (CCS) laid the foundations for the rigorous treatment of synchronisation and communication between processes, fundamental to the development of the distributed memory computation model of parallel computation.

Other notable highlights in this area have been the work in languages and applications. Not only has the design and implementation of individual languages led to a better understanding of the area of parallel computing but UK computer scientists have been at the forefront of international efforts to establish parallel software standards and benchmarks. Three notable examples are: MPI, the message-passing interface which is now widely accepted among vendors and researchers as a basis for promoting portability between high performance machines; High Performance Fortran Forum, which led to the design and development of a version of Fortran targeted to high performance machines; and Parkbench, a more scientific approach to benchmarking parallel machines.

Most recently UK computer scientists are playing a leading role in Java Grande developments. At the same time, much experimentation was carried out in the development of new compiler techniques to help with the efficient implementation of the linguistic concepts and the development of application software for these machines. To promote UK computer science expertise in this area, Parallel Application Centres were established at Queen Mary College, Oxford, Southampton and Edinburgh Universities. Not only did these Centres address a wide variety of application areas of benefit to other branches of science, they also collaborated with industry to demonstrate the commercial viability of parallel systems. Two of these Centres were designated as European 'Technology Transfer Nodes' and played a leading role in many EU projects bringing together applied parallel computing and industry. Applications range from engineering simulations with Computational Fluid Dynamics (CFD) or Finite Element Methods (FEM) to commercial applications such as data warehousing, data mining and financial modelling. Experience gained from these applications has been accompanied by research into software tools and parallel algorithms. For example, NASA currently funds development of CAP-Tools at Greenwich, a state-of-the-art tool for parallelising FEM codes.

Research on parallel algorithms ranges from work based on PRAMs and Valiant's BSP model to more pragmatic research on parallel numerical algorithms. NAG, for example, has produced a parallel version of their numerical libraries. There are also a number of recent spin-off companies developing and offering products to support cost-effective cluster computing.

Future computer science challenges are intimately related to the e-Science programme and the development of a reliable computation and information Grid infrastructure. Research into tools, algorithms and environments will all make significant contributions to the development of Grid middleware software.

Graphics and digital media

In the UK, Computer Graphics and its allied topics have a long history of blending practical application with underlying theoretical developments. As a result, much of its progress has directly involved user communities, including industry, with the result that there is no hard and fast division between university laboratories, government laboratories and the application developers. For example, as a result of university work in graphics and CAD in the 1970s, there was a major flowering of CAD companies and animation companies around Cambridge in the seventies and eighties.

The UK graphics community have been active in developing the standards GKS, CGM, PHIGS and CGI and participated extensively in refining the original Virtual Reality Modelling Language (VRML) specification for promulgation as an ISO/IEC standard.

Computer Graphics are resource-hungry, so there has been a continuing interest in hardware support for fast rendering. Work at Sussex in the 1970s (Grimsdale, Willis) established a group for rendering 3D scenes in real-time. Willis moved to Bath and his later work lead to the full-colour raster operation chip for the Atari Transputer Workstation, for which the same group wrote the graphics library. Meanwhile, the Sussex team played a major role in the SPIRIT Workstation (around 1990), which was produced within an ESPRIT Project to exploit new technologies in the field of computer graphics and other computationally intense tasks. The main partners in the consortium were University of Sussex (Grimsdale, Lister), University of Tuebingen (Strasser), Queen Mary and Westfield College (Slater, now at UCL), and the companies Caption (France) and Kontron (Germany). The resulting high performance Graphics System could attain real-time Gouraud shading at a rate of 300,000 triangles per second or 1,000,000 colour shaded vectors per second.

Virtual reality is now mature enough that it needs a deeper understanding of the principles of its application. The work of Slater (University College London) is notable here, with a series of experimental results achieved under carefully controlled conditions. His work on VR interaction between real users and virtual ones, the sense of immersion and what makes VR effective has been rewarded with a five-year Senior Research Fellowship from EPSRC. Recent developments have made complex real-time virtual realities possible. The team at Manchester (Hubbold) is the first both to understand the need to integrate a VR kernel into networks of workstations and to provide a solid implementation (Maverik: made available over the web). This is tested with exceptionally complex models provided by industry. The systems architecture approach both echoes the traditional Manchester strengths (which strongly influenced three generations of commercial computers) and moves forward to address modern needs.

Artificial intelligence and cognitive science

The UK has been a world-leader in Research in Artificial Intelligence and Cognitive Science (which for present purposes we take together to include Computer Vision, Neural Networks, and Computational Neuroscience) since the birth of the subject in the sixties, giving a home to the first departments in Europe in both areas, and hosting the 2nd International Joint Conference (IJCAI) in London in 1971 after the 1st in Washington D.C. in 1969. The UK has strengths and achievements in all the major areas of AI, and this is reflected by the fact that 7 out of approximately 50 members of the AI Journal editorial board are based in the UK (equal to the sum of the board members from Germany, France and Japan put together).

The UK was at the forefront of work in expert systems, producing the first commercial expert system shell, Expertise, which stood out from other shells of the time in its use of machine learning in the form of rule-induction. This, and several other large-scale expert system shells developed in the UK (e.g., XiPlus, RuleMaster and XTran), were deployed in a wide range of industrial applications including: chemical process control (Westinghouse), financial decision making (American Express UK), separation of gas from oil (BP), breakdown prevention of electrical transformers (Hartford Steam Boiler Company), configuration of building fire-detection equipment (Seimens), space shuttle engine testing (NASA), and forecasting severe thunderstorms (US National Weather Service).

The UK was, and still is, at the forefront of research in functional and logic programming, starting with early work in Edinburgh on the Pop2 (Popplestone, Burstall and others) and Prolog languages (Kowalski, Warren, Mellish). Many of the best known logic programming theoreticians (e.g., Clark, Imperial; Lloyd, Bristol) have also been UK-based. Poplog (Sloman at Sussex, now Birmingham), a multi-language AI development toolkit based originally on Pop2 and Prolog has been marketed commercially since 1982.

The UK has a long tradition of world-class research in the area of machine learning, beginning with Alan Turing's early investigations, Michie and Chambers' work on learning to balance a pole on a cart, and research by Michie, Niblett, Shapiro on Structured Induction. During the last decade, Inductive Logic Programming has become a widely recognised area of Machine Learning, which has its origins in the UK in early work by Plotkin (Edinburgh) and, more recently, Muggleton (York, Imperial). Valiant's (ex-UK now Harvard) research formed the basis of computational learning theory. In artificial neural networks, Willshaw and Longuet Higgins (Edinburgh) played a crucial role, and Hinton (now UCL) went on to become one of most influential figures in ANNs. Bishop's (Aston, now Microsoft Cambridge) contribution to the statistical foundations of the field are very important and influential. Commercial successes include Integral Solutions Ltd's Clementine software, which became market leader in machine learning for data mining. It incorporates neural net and rule induction packages with origins in research tools from Aberdeen and Sussex, and is by far the most commercially successful application of Poplog (Sloman). Clementine developers are now incorporating the UK-developed Inductive Logic Programming system, Progol, into the next release. (ISL was acquired by USA's SPSS Corporation in 1998.)

The UK is among the world leaders in computational linguistics and speech engineering, with many highly visible projects competing favourably in the international arena. The Cambridge HTK system, from Young's group, is consistently among the best on DARPA benchmark tasks, including: best system on Broadcast News 1997, joint equal best system on Broadcast News 1998, best system on Switchboard and CallHome 2000. HTK and other technologies from Cambridge formed the basis of Entropic Ltd, recently acquired by Microsoft. Edinburgh's Festival speech synthesis system, developed by Isard, Taylor, and Black is in use in hundreds of research labs and in industry around the world, including AT&T and Aculab, and forms the basis of a new spin-out company, Rhetorical Systems Ltd. Autonomy, founded by Lynch (Cambridge), exploits research on applying Bayesian techniques in natural language processing and data mining, and is now the world's leading provider of Internet based search technology, with a market capitalisation of £3B.

UK based work on planning began with Michie's (Edinburgh) early work on Graph Traverser, one of the earliest examples of heuristic search algorithms. The work of Tate and others at Edinburgh's AI Applications Institute on non-linear planning has consistently received

DARPA funding, and has been incorporated into numerous software products that have significant usage or commercial value, e.g., Formation, which is used to design the layout of all BT Yellow Pages; Expert Provisioner, which is saving some $30 million per annum for RAF Logistics and soon to be deployed to all forces; and OPlan, used in many NASA applications.

In theorem proving, Robinson's was central to most modern automated theorem proving. Early work by Meltzer (Edinburgh) and others was also very influential. The work of Bundy's (Edinburgh) DREAM group is widely recognised internationally. Cohn's (Leeds) work in spatial logics is also highly regarded, as is Frisch's (York) work on constraint solving.

In computational vision, David Marr and Richard Gregory (Bristol) have been extremely influential figures. More recently, Brady's (Oxford) work on medical image understanding leads the field. In addition, work by Blake (Microsoft Cambridge), Maybank (Reading) and Hancock (York) is highly regarded. Robotics projects in Edinburgh, e.g., Freddy (Michie, Barrow, Popplestone, Ambler) led the international scene in the late 1960s, demonstrating the potential for automated assembly. Although the funding situation has been difficult, robotics research teams, such as that of Brady (Oxford), are once more extremely strong.

Space prohibits listing many other excellent groups and researchers, but it is noteworthy that the UK has been able to attract a number of outstanding younger faculty with growing international reputations, including Gaizauskas (Sheffield), Dayan (UCL), Ghahramani (UCL), Husbands (Sussex), D. MacKay (Cambridge), Preese (Aberdeen), C. Williams (Edinburgh) and Jennings (Southampton, winner of the 1999 IJCAI Computers and Thought Award), among others.

Artificial Intelligence is the part of computer science that works at the boundaries of the subject and at the interface with other disciplines. (The tendency of its discoveries to be appropriated by other computer science subdisciplines is a direct result). Future research in AI will undoubtedly continue to have this character, fusing previously disparate fields and collaborating across disciplinary boundaries. Among the grand challenges that currently face all subfields of Artificial Intelligence is the unification of probabilistic and machine learning-based methods with high level symbolic representation and reasoning.

Human computer interaction

HCI has a pivotal role in computing to ensure that computers are usable, effective and efficient. There has been a long tradition of HCI research in the UK beginning after the Second World War and concertedly from the late 1960’s with the work at the MRC Applied Psychology Unit Cambridge. From a psychological perspective research from that Unit (Morton, Barnard, Long, and Young) formed an embryonic core focusing upon cognitive modelling.

A major UK strength in HCI is in investigating the use of computers to support people’s work. The UK leads research employing multidisciplinary perspectives to analyse, describe, model and understand interaction behaviours. Research in cognitive modelling in HCI is currently especially strong at Cambridge (Barnard), Bath (H. Johnson, Duke), UCL (Long), Hertfordshire (Young) and Sheffield (May). This is applied to both design and more especially evaluation methods, Bevan (NPL), Harrison and Monk (York), Johnson (Glasgow) and Long (UCL), the latter leading to the MUSE methodology. Robinson (Cambridge) develops new modes of interaction in his active desk.

An important area of research in the UK is HCI for the disabled, where Newell and colleagues at Dundee lead the world in this small but very important area of research. A further strength is in carrying out empirical investigations of novel and alternative forms of interaction, for example, Edwards (York), Brewster (Glasgow) work on haptic interfaces for the visually impaired. Advances in information visualisation have been achieved by Spence (ICSTM) and Chalmers (Glasgow).

The sociological aspects of HCI are also strong in the UK: Suchman (Lancaster) and Heath (Kings College London) and in UK-based industrial research laboratories Xerox (Cambridge) and HP Labs (Bristol). Researchers at Lancaster and Nottingham (Rodden) work in the area of computer-supported co-operative working and “virtual environments” where group working and social interaction form an important key to the development of distributed and virtual reality systems.

The UK is also especially strong in the development of theoretical and methodological approaches to HCI design: Long (UCL), Johnson (Bath), Sutcliffe (UMIST). This has led to the establishment of domain, work and user-interface modelling in integrated HCI and software engineering methods, the development of model-based user interface design methods and tools, and requirements analysis. Harrison (York), Duke (Bath), and Johnson (Glasgow) work in formal methods in HCI. Interesting and leading developments in this area include interactor modelling, Harrison (York), syndetic models of cognitive and computational systems (Bath) and the application of formal methods to safety-critical systems C. Johnson (Glasgow and York).

Research at the Royal College of Art (Crampton-Smith) in the area of graphic design provides an important aesthetic contribution to user-interface design: Edmonds (Loughborough) works on creativity and computing and music is strong at Bath, Loughborough, Edinburgh and Glasgow.

HCI is now a major factor in computing and related industries, especially in web/internet design, service provision and games, and numerous industrial research labs in the UK and overseas employ HCI practitioners. It is increasingly important in dependability: a future consideration will be the ability to consider emotional factors to accommodate such attributes as enjoyment, anxiety and stress as important factors in user interface design and evaluation. For example, Draper (Glasgow) is developing an understanding of fun, for the games industry.

Information engineering

Information engineering supports all of the processes necessary for an information-powered organisation or society. These include information creation or capture, information description and storage, the integration of information and the extraction of relevant subsets and their use to inform people or to derive further information. Requirements, practices and expectations vary rapidly and the available technology is continually advancing. As a result, flexibly supporting change and heterogeneity is a fundamental requirement that will always be important.

The UK has several groups that contribute significantly to internationally renowned research in these areas. Key contributions in databases include early work on OODB, languages and environments: Atkinson (Glasgow) and Morrison (St Andrews), active databases: Paton (Heriot-Watt and Manchester), and parallel databases Watson (Newcastle). Other noteworthy work includes transactions in databases Moody and Bacon (Cambridge), logic and algebra for databases P. Gray (Aberdeen), constraints A. Gray (Cardiff), theory Levene (UCL), and deductive systems: Poulovassilis (Birkbeck) and Williams (Heriot-Watt).

The technology to support advances in information systems includes the integration of programming languages with databases, simplifying the construction of higher-level systems: Atkinson and Printezis (Glasgow), Morrison, Dearle and Kirby (St Andrews); leading to innovations in the Java platform to support safe dynamic schema, program and data evolution. The Arjuna project, Shrivastiva (Newcastle), the work of Bacon and Moody (Cambridge), the

Commandos project: Harper (Glasgow, RG), Norrie (Glasgow, ETH), and Baker (Dublin, Iona) and the ANSA project, Herbert (Cambridge) have delivered distributed object stores and heterogeneous interworking, that has led to products and standards. Connor (Strathclyde) and Cardelli (Microsoft Research, Cambridge) develop languages and type systems to support network-based heterogeneous computation. Software architectures supporting flexible integration have been developed by Welland (Glasgow), and by Blair, Hutchinson, Rashid (Lancaster), the latter particularly for multi-media applications. These approaches begin to exploit extended metadata and reflection. The general thrust of attempting to improve the description of data and exploit these descriptions in information integration is helped by improving standard infrastructure, such as Thompson's contribution (Edinburgh) to XMLSchema.

The UK can claim a significant world lead in areas of information retrieval. Van Rijsbergen's (Glasgow) 'Logical Model', pioneered in 1986, spawned a complete new subfield of IR, and a commercial product HySpirit. With Spark-Jones (Cambridge) he laid down the initial design for the TREC evaluation methodology. His original book is the most cited IR textbook and he was the first researcher to claim effectiveness gains for clustering. Spark-Jones developed the probabilistic model for IR in the seventies and continues to lead work on NLP/IR. The widely used commercial search engine, Muscat, grew out of the joint work of Porter, Robertson and Van Rijsbergen in Cambridge. Other notable IR research centres include Sheffield (Willett, Sanderson, Gaisauskas), Strathclyde (Crestani, Dunlop) and QMW (Lalmas).

Knowledge Management covers many different research topics. Key is the area of ontologies in which the UK is one of the world leaders, notably Manchester (Goble, Horricks), and OU (Motta). In particular the team at Manchester has led the way for biomedical systems and healthcare, and has led the development of the OIL language for the Web. Southampton (Hall) are world leaders in the development of hypermedia and web technology. In particular their Microcosm system is cited as one of the key influences on XLink, and their dynamic link service work has been successfully commercialised. Other key players in the UK development of Web technology include Bristol (RDF), Edinburgh (XMLSchema/RDF) and the OU.

Future challenges include the evolution of database technology to cope with very large repositories of unstructured or semi-structured information in order to work with the web, and the integration of database systems, IR and KBS. One of the main drivers here is the concept of the "semantic web" and the increasing use of metadata and ontologies to facilitate the storage and retrieval of information. Major advances are required in search engine technology but this will be complemented by the application of advanced hypermedia technology. One of the recently funded Interdisciplinary Research Collaborations (IRC) is in Advanced Knowledge Technology (AKT) which seeks to integrate work on the six components of the knowledge life cycle, namely acquisition, modelling, reuse, retrieval, publication and maintenance. The partners are Southampton, Aberdeen, Edinburgh, Sheffield and the Open University.

Software engineering

The UK has made substantial contributions to software engineering from the outset, with Randell (Newcastle) and Buxton (then, Warwick) being editors of the original NATO conference proceedings where the term was coined. Contributions are mainly in the areas of requirements engineering, software processes, system architecture and evolution, real-time systems, testing and analysis.

Centres for research in system requirements engineering and system evolution include Durham, Imperial, Lancaster, UMIST, York and Southampton. Some of the more novel results in these areas come from collaboration with social scientists concerning the engineering process (Sommerville, Lancaster); systems dynamics for modelling system evolution (Lehman, Imperial), and practical tools for legacy system re-engineering which have been used in industry (Bennett, Durham).

Real-time systems research includes priority-based real-time system scheduling, and static worst-case execution time methods and tools, many of which are now being used in commercial products. For example, in long term collaborative projects with British Aerospace, McDermid, Wellings and Burns (York) have made significant impact on real time systems development for aerospace and automotive embedded systems. System architecture and process results include support for dynamic system configuration giving the ability to change parts of the system while the rest remains operational (Kramer, Magee, Imperial), techniques for family analysis and design of embedded systems (McDermid, York), and process modelling languages for process automation (Warboys, Manchester).

Software testing and metrics are a strength at Liverpool, Brunel, Keele, Sheffield, and York, focussing on metrics for control-flow based test coverage (Hennell, Liverpool); test automation, techniques and coverage metrics for OO software (Clarke, McDermid, York); principles of software product and process metrics (Fenton, City; Kitchenham, Keele). Carre (Southampton) pioneered work in static analysis leading to commercial static program analysis tools and program verification environments.

These research advances have also influenced software engineering education, with major textbooks and reference books being produced by McDermid and Sommerville. Sommerville’s text (first produced in 1982 and now in its sixth edition) is internationally recognised as an essential undergraduate text: it has 1/3 of the international market, is adopted in over 400 institutions, and more than 30,000 copies will be sold this year.

The key role of software engineering research will be to continue to develop new techniques and technologies to support software development. There is no question that software engineering has matured and, in spite of the problems of software projects, we do now routinely build large, complex, reliable software systems. However, new challenges face the research community including: supporting accelerated development processes without compromising the quality of the system; developing evaluation frameworks for new methods and techniques based on a deep understanding of relationships between organisations and technology and ways of relating small-scale systems to large-scale systems etc., and discovering effective ways of using new technology such as mobile systems, large screens, etc. to support software engineering processes such as outsourcing, distributed negotiation, etc.

Dependability

Techniques for achieving dependability in the presence of intermittent hardware failure or decay are well understood; there remains a great need for research in the creation of systems that are dependable despite residual design errors and/or external interference. Progress on this research is essential if complex hardware/software systems are to be made dependable for use in safety-critical and business-critical applications.

The UK was one of the first in this area; Brian Randell at Newcastle played a central role in establishing the field and many key figures in the US today were in Newcastle (Rushby, Cambell, Mellior-Smith, Horning, Lomet, etc.). The Newcastle team devised a seminal architectural concept for coping with design errors: “recovery blocks” provide a way of making alternative algorithms available at run time. This strand of research has led more recently to an extension of the concept of transactions called Coordinated Atomic Actions which has been embodied in products from Arjuna Solutions Ltd.

The UK is also strong in security and cryptography. Many of the applications of dependability (particularly security) research are secret or commercially confidential; indeed it only recently became widely known that public key encryption was devised by UK government security researchers around 1970 (prior to the usually cited RSA inventors). Currently there are strong groups at Cambridge and RHUL, at the government security labs in Cheltenham and at Microsoft Research. Ross Anderson at Cambridge is a noted international expert on technology and policy matters, and created Serpent, which recently came second in the competition to find AES, a replacement for DES. Piper and Mitchell head the group at RHUL which spawned the security company Zergo: recently taken over by Irish Company Baltimore. Important research on ways of reasoning about security include Cambridge's "BAN logic" and subsequent developments ("Ambients") in Microsoft's Cambridge Research Lab. by Cardelli and Gordon. Another key outcome is the use in Oxford of the "FDR" tool to analyse the safety of protocols.

In the area of software reliability, UK researchers have developed and validated theories for modelling software reliability growth which have been used to assess critical systems, and to influence standards for assessing critical software (Littlewood, City). In research on Safety Cases, the UK has pioneered techniques for structuring safety arguments and safety cases which are now supported by commercial tools, and used in a variety of industries, e.g. aerospace, naval nuclear power, and railways (McDermid, Kelly, York);

In spite of the inherent secrecy of some dependability research, there is an excellent public record of industrial impact over many years, e.g. York and Newcastle running the BAe-sponsored Dependable Computing Systems Centre and the Safety Critical Systems Club run by Newcastle and City.

One major research challenge is for techniques to build dependable systems from pre-existing (possibly COTS) systems that are not themselves dependable; these are issues that the EPSRC DOTS project and the, UK-led, EU-funded DSoS project are tackling. Malicious attacks are an ever-growing problem and are addressed in the, UK-led, EU-funded project MAFTIA. Designing dependable systems that include humans poses a huge interdisciplinary research challenge that is being tackled by the five-university IRC on the Dependability of Computer-Based Systems.

More research is also needed in logics and notations for reasoning about dependability; this requires both stochastic reasoning and arguments about belief, trust and knowledge.

Programming languages

The UK has a strong tradition of influential work in the design and implementation of programming languages. Christopher Strachey pioneered key concepts in language design such as polymorphism and the orthogonality of language features; his CPL project led to BCPL (Martin Richards) and thence (via B) to C.

The UK has also played a leading role in the development of declarative programming paradigms and languages. In functional programming, Standard ML [Milner et al. Edinburgh] has been enormously influential in its pioneering use of types, combining the discipline of strong typing with the flexibility of polymorphic type inference. Turner [Kent] pioneered the lazy functional paradigm with his language Miranda, while Peyton-Jones [Glasgow, now Microsoft Cambridge] is a leading figure in the development of Haskell. There has been an equally strong UK role in the development of logic programming, with leading figures such as Kowalski [Imperial], one of the founders of the paradigm, Warren [Bristol] who designed its standard implementation and went on to found Quintus, Clarke [Imperial] who developed PARLOG (and whose student Foster, now at Argonne, was one of the leading figures in the development of the USA computational grid) and Lloyd [then Bristol], who has developed the languages Godel and Escher.

There has also been an important contribution to the design of languages for concurrent programming, notably Occam [May, Bristol] based on Hoare's theory of Communicating Sequential Processes. Hoare's earlier work on monitors as a language construct for shared-variable concurrency, and on data abstraction, was a key influence on the design of Java. Milner's CCS provided the basis for LOTOS, the ISO specification language for communications protocols. The pi-calculus [Milner; Parrow; Walker, Oxford] is also proving seminal for a number of current languages and formalisms aimed at providing a sound basis for today's 'mobile' and 'global' computation, notably the work on Ambients by Cardelli and Gordon [Microsoft, Cambridge], PICT [Pierce and Turner], Distributed Pict [Sewell, Cambridge] and the join calculus [Levy, Gonthier, Fournier]. The Spi calculus [Abadi and Gordon] has been applied to security, while the Stochastic Pi-calculus [Shapiro] has been used to model signal transduction in biomolecular processes. It is noticeable that Microsoft's new Cambridge research lab has hired many programming languages experts, and commercial exploitation of ideas of include proof carrying code and parametric polymorphism applied to Java, and the use of formal methods in verifying the Mondex smart card, carried out at Logica.

Bioinformatics

Bioinformatics is the application of Computer Science theory and technology to support research and development in the Biological Sciences. It covers the acquisition, management and analysis of biological information in all its forms. The advent of rich supplies of data from laboratory mechanisation in biology and the acceleration of research progress through web-based data exchange is stimulating a demand for innovations from the computing community. Study of biological systems is also suggesting new approaches to computational challenges.

Bioinformatics applies existing techniques developed in Computer Science to problems in Biology; for example information management or high performance computing. Models in molecular biology are becoming so complicated that computers are essential. To build these models, and to support the collaboration, information and knowledge these models depend upon, needs new computer science. New levels of computational power are demanded to model the functional behaviour of cells. New levels of integration are demanded to handle the range of scientists and techniques that interact here, from the biochemists, molecular biologists, geneticists and crystallographers to engineers, chemists, physicists, pharmacologists, ecologists, epidemiologists, medical and veterinary researchers. This challenge will be a major interest in the next decades and responding to it will extend the capabilities of computing science in ways applicable in many other domains.

The UK Computer Science and Bioinformatics interface has flourished under the EPSRC-BBSRC Bioinformatics Initiative. In 2000 there were 70 applications to the final round of the initiative, and excellent proposals were made. Every major CS department now has an activity in Bioinformatics. The EPSRC Life Sciences Interface initiative has funded 4 chairs in Bioinformatics at Manchester, UCL, Imperial and Oxford. The UK is currently a world leader in collaborative bioinformatics. Bioinformatics and functional genomics appear as a key driver for the OST's E-Science initiative.

The UK has the European Bioinformatics Institute (EBI) at Hinxton, co-located with the Sanger Centre on the Wellcome Genome Campus. The EBI is the custodian and distributor of the major information repositories and tool sets used throughout the world by biologists. The EBI relocated to the UK from Germany in 1994 through support by the Wellcome Trust and the MRC. The EBI has a healthy and active collaboration with UK academic bioinformaticians.

The UK is responsible for the publication and curation of major databases such as EMBL (EBI) PRINTS (Attwood at Manchester) and CATH (Orengo and Thornton at UCL), database systems such as AceDb (Durbin at the Sanger Centre) and MaxD (Brass and Rattery at Manchester), tools such as GASP (Willett at Sheffield) and standards such as ArrayExpress (Robinson at EBI).

Willett (Sheffield) and Hunt (Glasgow) have developed novel algorithms, data structures and tools for representing, searching and processing the information in biological and chemical databases.

Pharmaceutical companies, and biotechs, have traditionally had close research relationships with the biological sciences research departments. This has continues in Bioinformatics, where a considerable amount of the funding comes from industry; for example the TAMBIS, MaxD, GIMS and CINEMA projects at Manchester were funded/supported by AstraZeneca, Aventis, GlaxoWellcome and Pfizer respectively; GASP at Sheffield funded by Tripos Inc; ILP machine learning supported by SmithKlineBeecham and GlaxoWellcome.

Theory, formal methods and foundations

The UK has always been strong in semantics and foundations. Hoare and Milner both hold Turing awards; Hoare, Milner, Plotkin and Gordon are Fellows of the Royal Society and Hoare received the Kyoto prize in 2000. Strong research achievements past and recent have been in programming logics [Hoare, Oxford], operational semantics and domain theory [Plotkin, Edinburgh], models of concurrency [Hoare; Milner, Cambridge], game-theoretic semantics [Abramsky, Ong, Oxford; Hyland, Cambridge], exact real computation [Edalat, ICSTM].

This work is strongly linked to the development of methodologies and tools. The former include VDM (initiated in Vienna but progressed by Jones on his return to Manchester) and Z (initiated by Abrial, developed in his stay in Oxford, further refined by Sufrin et al. after Abrial returned to France) and the latter include automated reasoning systems such as HOL [Gordon, Cambridge], Isabelle [Paulson, Cambridge], FDR [Roscoe, Oxford] and the Edinburgh Concurrency Workbench; these have been used in the verification of both hardware (for example by Intel for their latest I-64 chip) and software. UK research is playing a leading role in the currently very active application of formal methods to security protocols [Roscoe; Lowe, Oxford; Paulson; A. Gordon, Microsoft].

Theory has informed strong programming language research [Cambridge, ICSTM, Manchester, Edinburgh, Oxford] which is the topic of another section. The PEPA Workbench [Hillston, Edinburgh] has pioneered probabilistic process algebra and its application to performance analysis; there is active work on probabilistic model-checking [Kwiatkowska, Birmingham] and on formal methods for probabilistic algorithms [Morgan and Sanders, Oxford].

The UK is also strong (though in smaller volume) in algorithms and complexity [Paterson, Warwick; Jerrum, Edinburgh; Brent, Oxford] and computer algebra [Davenport, Bath]. Jerrum was awarded the Godel Prize for his work with Sinclair on rapidly mixing Markov chains. Recent work has found applications in software for parallel algorithms [McColl, Oxford], novel approaches to computational mathematics [Martin, St Andrews] and bioinformatics [Paterson, Warwick]. Interdisciplinary theoretical work is supported in particular by the EPSRC/LMS MathFit initiative.

Computing has grown from stand-alone machines and programs, under designer control, to a mobile worldwide system of interacting hardware and software that is designed in the small but is haphazardly assembled. A tremendous challenge is to build mathematical models of this assembly, which via scientific consensus will bring robust qualities to the global network.

There are many parts to this challenge: logical and process-theoretic models in which to analyse what can and can’t happen; stochastic and quantitative models in which to analyse throughput efficiency against many different measures; fundamental semantic models, built upon game theory and allied theories, integrating programming at the small scale with massive combinatorial structures at the large; formal techniques, languages and allied tools for mediating these models to designers and clients.

B The e-science opportunity

The increase in simulation and modelling, the collection and storage of large databases, the need for varied computational facilities and the collaboration between personnel across locations underlines the challenge of e-science. This is apparent in a wide range of scientific and engineering applications and has substantial benefits for commercial and industrial applications. The underlying concept on which e-science is based is the Grid, a high-speed network that links computers, specialised processors, instruments, and databases. It differs from existing technologies mainly through its middleware that make collaborative computing easier and more reliable. The objective is to provide the on-demand creation of powerful virtual computing systems access to all resources. It is particularly appropriate now because of increasing bandwidth and services that are available, the advances in storage capacity, the increased availability of compute resources and the advances in application concepts.

The grid requires new tools and programming methods based on a rethink of the entire program development cycle with run-time adaptation as a basic concept. This calls for an interdisciplinary research agenda: many areas in which UK computer science is strong match the development and implementation of grid technologies and their use in e-science.

Structure The grid can be regarded as a facility with a number of layered components, some of which are well understood and others which, at the moment, are only speculative. This gives rise to a large number of computer science research issues. The first layer is the computational grid that provides the raw computing power, high bandwidth interconnection and associated data storage. The second layer, the information grid, will allow easily accessible connections to sources of information and tools for its analysis and visualisation. The next layer, the knowledge grid, involves techniques such as data mining and machine learning that give added value to the information and also provide intelligent guidance for decision-makers. As more understanding emerges, other layers may develop.

More specifically, the requirements for the computational grid are tools and systems to handle performance, resilience, distribution of computation and data, resource balancing, scalability, locality of compute power and data resources, authenticity and authorisation, multiple utilisation of the infrastructure and performance. The information grid should provide tools and systems for homogeneous access to heterogeneous representations, location of information, information/content distribution, inter process communication, active content management, multicast issues, middleware for information access and integration, longer term scheduling, consistency management version. There is an urgent requirement for tools and systems to support standards and meta-data, data mining, knowledge, elicitation and support for hypothesis and credibility measures. Systems to support e-literature and grey literature including extraction of encapsulated knowledge. Overriding all three layers is the need for a user workbench where users interact with the grid's facilities.

Research Issues Long term the development and deployment of ubiquitous, seamless, transparent grid technologies is needed. In addition, research and innovation are required in the development of advanced applications such as coupled simulations, self-scheduling computations and the solution of complex meta-problems. Some of the research issues that need to be addressed include: Heterogeneity, scalability and dynamicism of hardware and software resources; System installation, configuration and management; Remote data access and movement; Resource management and scheduling; Performance modelling and evaluation; security and fault tolerance; interoperability/extensibility and the interaction with mobile and other devices

C Barriers to exploiting the science base

The OST report, Technology Matters [Ref-8], provides an extensive analysis of issues surrounding UK exploitation of the science base:

- The UK's under performance in creating wealth from its world class science base has been the subject of concern, inquiry and study ever since the country lost its industrial and technological leadership towards the end of the 19th century. Successive Governments over the years have pursued an extensive range of different priorities and approaches. Yet the nature of the problem, possible solutions and the role of the Government remain among the most vexed policy issues of the age.
- Among the commonly cited reasons are an anti-technology culture in the UK; a failure or inability within companies to exploit the opportunities of technology; more interest in science or financial performance than technology based innovation; a lack of enterprise, ambition, vision, drive and entrepreneurship; management and managerial weaknesses; risk aversion and too great a focus on the short term.
- Others include a relatively small, undemanding domestic market; frequent boom and bust cycles with widely fluctuating interest and exchange rates; the haemorrhaging of the manufacturing base during the 1980s with two recessions and extensive corporate raiding by so called asset strippers; failings in Government policies, and a widely perceived flight of the nation's most able young people from careers in science, engineering and technology in favour of such professions as medicine, finance, law, accountancy and the media.

They recommend in particular:

- broadening and deepening the UK's technology base, particularly in the strategically important sectors of the 'sun-rise' industries, such as computing, microelectronics, telecommunications, advanced materials and bio-technology;
- developing more dynamic clusters and strengthen links to the increasingly globalised sources of leading edge S&T worldwide and
- strengthening the nation' s ability to create new technology based businesses and to grow them rapidly into large companies, in particular through financial resources
- These further policy developments should be directed at meeting the need for:
  - government, and society more widely, to place technology on a par with science in terms of the standing, recognition and value attached to each of them;
  - a larger cadre of people who are highly skilled both as active technologists and in running a business;
  - personal finance and tax incentive structures to be fully competitive internationally, especially with those in the US: business and the providers of financial resources need to become fully alive, and responsive, to the technology challenge; and
  - the Government to play its full part in driving exploitation of the new technologies in all ways it interacts with business - sponsorship in the widest sense.

D Membership of UKCRC

(Affiliations shown only for industrial members.)

Convenor: Professor Alan Bundy, FRSE, FRSA, FAAAI - A.Bundy@ed.ac.uk

Professor Malcolm Atkinson, FRSE, FBCS - mpa@dcs.gla.ac.uk

Professor Richard Bird - richard.bird@comlab.ox.ac.uk

Professor Muffy Calder - muffy@dcs.gla.ac.uk

Mr Brian Collins, (International Director of IT, Clifford Chance) - Brian.Collins@cliffordchance.com

Professor Michael Denham - M.Denham@plymouth.ac.uk

Professor Steve Furber FREng, FBCS - sfurber@cs.man.ac.uk

Professor Wendy Hall CBE, FREng, FBCS - W.Hall@ecs.soton.ac.uk

Professor Chris Hankin - clh@doc.ic.ac.uk

Professor Tony Hey, FIEE, FBCS - ajgh@ecs.soton.ac.uk

Professor Sir Tony Hoare, FRS, FBCS, M.Acad.Eur (Microsoft Research Ltd) - thoare@microsoft.com

Professor Andy Hopper FREng, FIEE (AT&T Laboratories) - hopper@eng.cam.ac.uk

Professor Cliff Jones, FIEE, FBCS, FACM - Cliff.Jones@ncl.ac.uk

Professor Derek Mcauley (Marconi Research) - Derek.Mcauley@marconi.com

Professor John McDermid EurIng, FIEE, FRAeS, FBCS - john.mcdermid@cs.york.ac.uk

Professor David May FRS - dave@cs.bris.ac.uk

Professor Ursula Martin - um@dcs.st-and.ac.uk

Professor Robin Milner FRS, FRSE, Dist FBCS, M.Acad.Eur - rm135@cl.cam.ac.uk

Professor Johanna Moore - J.Moore@ed.ac.uk

Professor Roger Needham FRS, FREng, FBCS, M.Acad.Eur (Microsoft Research Ltd.) - needham@microsoft.com

Professor Mike Paterson FRS, M.Acad.Eur - msp@dcs.warwick.ac.uk

Professor Ron Perrott FBCS, FRSA, FACM - r.perrott@qub.ac.uk

Professor Ian Sommerville FIEE - is@comp.lancs.ac.uk

Professor Martyn Thomas FIEE, FBCS - martyn@thomas-associates.co.uk

Professor Ian Wand FIEE - icw@cs.york.ac.uk

Professor Ian Watson - watson@cs.man.ac.uk

Professor Phil Willis - P.J.Willis@bath.ac.uk

Secretary Mrs Peta Walmisley, External Relations Manager, The British Computer Society - pwalmisley@hq.bcs.org.uk

E Bibliography of supporting materials

(The material with bracketed references available at www.dcs.st-and.ac.uk/~um/ukcrcir.)

[Ref-1] INFORMATION AND COMMUNICATION TECHNOLOGIES IN THE UK, A SCENARIO FOR SUCCESS IN 2005. *(UKCRC1_ICT2005.pdf)*

[Ref-2] Excellence and Opportunity a science and innovation policy for the 21st century, DTI While Paper 3000. *(UKCRC2_whitepaperwhole.pdf)*

[Ref-3] Opportunities for realising our potential excellence, Save British Science paper 01/01. *(UKCRC3_SBS0101.pdf)*

[Ref-4] Computer Science: the Core IT Discipline, Robin Milner and others 1991. *(UKCRC4_milner1991.pdf)*

[Ref-5] EPSRC IT Sector Evaluation Programme Report, NCC August 2000. *(UKCRC5_NCCFinalReport.pdf)*

[Ref-6] Fundamentals of Computer Science, Revised Project Description, National Academy of Science CSTB. *(UKCRC6_NASCSTB.pdf)*

[Ref-7] The Impact Project: Determining the Impact of Software Engineering Research upon Practice, NSF. *(UKCRC7_NSFSEImpact.pdf)*

[Ref-8] TECHNOLOGY MATTERS, Report on the exploitation of science and technology by UK business , OST February 2000. *(UKCRC8_TechMatters.pdf)*

[Ref-9] Software Technologies Review Recommended EPSRC classification of computer science research EPSRC classification of computer science research EPSRC. *(UKCRC9_SoftTecRev.pdf)*

[Ref-10] PITAC report, Executive summary. *(UKCRC10_PitacExecSum.pdf)*

[Ref-11] EPSRC classification of computer science research. *(UKCRC11_ITCSkey)*

[Ref-12] Computing Tomorrow. (eds.) Ian Wand and Robin Milner, Cambridge University Press, 1996.

[Ref-13] Informatics: 10 Years back, 10 Years Forward, (ed.) Reinhard Wilhelm, Springer-Veralg, LNCS 2000, 2001.

[Ref-14] The Arjuna story *(UKCRC14_Arjuna.pdf)*